

HW4

February 19, 2020

1 CSE 252B: Computer Vision II, Winter 2019 – Assignment 4

1.0.1 Instructor: Ben Ochoa

1.0.2 Due: Wednesday, March 4, 2020, 11:59 PM

1.1 Instructions


- Review the academic integrity and collaboration policies on the course website.
- This assignment must be completed individually.
- All solutions must be written in this notebook
- Math problems must be done in Markdown/LATEX.
- You must show your work and describe your solution.
- Programming aspects of this assignment must be completed using Python in this notebook.
- This notebook contains skeleton code, which should not be modified (this is important for standardization to facilitate efficient grading).
- You may use python packages for basic linear algebra, but you may not use packages that directly solve the problem. If you are uncertain about using a specific package, then please ask the instructional staff whether or not it is allowable.
- You must submit this notebook exported as a pdf. You must also submit this notebook as an .ipynb file.
- You must submit both files (.pdf and .ipynb) on Gradescope. You must mark each problem on Gradescope in the pdf.
- It is highly recommended that you begin working on this assignment early.

```
In [1]: %matplotlib inline
import numpy as np
from PIL import Image
import matplotlib.pyplot as plt
import matplotlib.patches as patches
import time

# open the input images
I1 = np.array(Image.open('price_center20.JPG'), dtype='float')/255.
I2 = np.array(Image.open('price_center21.JPG'), dtype='float')/255.

# Display the input images
plt.figure(figsize=(14,8))
plt.subplot(1,2,1)
```

```
plt.imshow(I1)
plt.subplot(1,2,2)
plt.imshow(I2)
plt.show()
```



1.2 Problem 1 (Programming): Feature detection (20 points)

Download input data from the course website. The file price_center20.JPG contains image 1 and the file price_center21.JPG contains image 2.

For each input image, calculate an image where each pixel value is the minor eigenvalue of the gradient matrix

$$N = \begin{bmatrix} \sum_w I_x^2 & \sum_w I_x I_y \\ \sum_w I_x I_y & \sum_w I_y^2 \end{bmatrix}$$

where w is the window about the pixel, and I_x and I_y are the gradient images in the x and y direction, respectively. Calculate the gradient images using the fivepoint central difference operator. Set resulting values that are below a specified threshold value to zero (hint: calculating the mean instead of the sum in N allows for adjusting the size of the window without changing the threshold value). Apply an operation that suppresses (sets to 0) local (i.e., about a window) non-maximum pixel values in the minor eigenvalue image. Vary these parameters such that around 600–650 features are detected in each image. For resulting nonzero pixel values, determine the subpixel feature coordinate using the Forstner corner point operator.

Report your final values for:

- the size of the feature detection window (i.e. the size of the window used to calculate the elements in the gradient matrix N)
- the minor eigenvalue threshold value
- the size of the local nonmaximum suppression window
- the resulting number of features detected (i.e. corners) in each image.

Display figures for:

- original images with detected features

A typical implementation takes around 30 seconds to run. If yours takes more than 60 seconds, you may lose points.

```
In [ ]: def ImageGradient(I):
    # inputs:
    # I is the input image (may be mxn for Grayscale or mnx3 for RGB)
    #
    # outputs:
    # Ix is the derivative of the magnitude of the image w.r.t. x
    # Iy is the derivative of the magnitude of the image w.r.t. y

    m, n = I.shape[:2]

    """your code here"""

    return Ix, Iy

def MinorEigenvalueImage(Ix, Iy, w):
    # Calculate the minor eigenvalue image J
    #
    # inputs:
    # Ix is the derivative of the magnitude of the image w.r.t. x
    # Iy is the derivative of the magnitude of the image w.r.t. y
    # w is the size of the window used to compute the gradient matrix N
    #
    # outputs:
    # J0 is the mnxn minor eigenvalue image of N before thresholding

    m, n = Ix.shape[:2]
    J0 = np.zeros((m,n))

    #Calculate your minor eigenvalue image J0.
    """your code here"""

    return J0

def NMS(J, w_nms):
    # Apply nonmaximum suppression to J using window w_nms
    #
    # inputs:
    # J is the minor eigenvalue image input image after thresholding
    # w_nms is the size of the local nonmaximum suppression window
    #
    # outputs:
```

```

# J2 is the max resulting image after applying nonmaximum suppression
#

J2 = J.copy()
"""your code here"""

return J2

def ForstnerCornerDetector(Ix, Iy, w, t, w_nms):
    # Calculate the minor eigenvalue image J
    # Threshold J
    # Run non-maxima suppression on the thresholded J
    # Gather the coordinates of the nonzero pixels in J
    # Then compute the sub pixel location of each point using the Forstner operator
    #
    # inputs:
    # Ix is the derivative of the magnitude of the image w.r.t. x
    # Iy is the derivative of the magnitude of the image w.r.t. y
    # w is the size of the window used to compute the gradient matrix N
    # t is the minor eigenvalue threshold
    # w_nms is the size of the local nonmaximum suppression window
    #
    # outputs:
    # C is the number of corners detected in each image
    # pts is the 2xC array of coordinates of subpixel accurate corners
    #     found using the Forstner corner detector
    # J0 is the max minor eigenvalue image of N before thresholding
    # J1 is the max minor eigenvalue image of N after thresholding
    # J2 is the max minor eigenvalue image of N after thresholding and NMS

    m, n = Ix.shape[:2]
    J0 = np.zeros((m,n))
    J1 = np.zeros((m,n))

    #Calculate your minor eigenvalue image J0 and its thresholded version J1.
    """your code here"""

    #Run non-maxima suppression on your thresholded minor eigenvalue image.
    J2 = NMS(J1, w_nms)

    #Detect corners.
    """your code here"""

    return C, pts, J0, J1, J2

```

```

# feature detection
def RunFeatureDetection(I, w, t, w_nms):
    Ix, Iy = ImageGradient(I)
    C, pts, J0, J1, J2 = ForstnerCornerDetector(Ix, Iy, w, t, w_nms)
    return C, pts, J0, J1, J2

In [ ]: # input images
I1 = np.array(Image.open('price_center20.JPG'), dtype='float')/255.
I2 = np.array(Image.open('price_center21.JPG'), dtype='float')/255.

# parameters to tune
w = 1
t = 1
w_nms = 1

tic = time.time()
# run feature detection algorithm on input images
C1, pts1, J1_0, J1_1, J1_2 = RunFeatureDetection(I1, w, t, w_nms)
C2, pts2, J2_0, J2_1, J2_2 = RunFeatureDetection(I2, w, t, w_nms)
toc = time.time() - tic

print('took %f secs'%toc)

# display results
plt.figure(figsize=(14,24))

# show pre-thresholded minor eigenvalue images
plt.subplot(3,2,1)
plt.imshow(J1_0, cmap='gray')
plt.title('pre-thresholded minor eigenvalue image')
plt.subplot(3,2,2)
plt.imshow(J2_0, cmap='gray')
plt.title('pre-thresholded minor eigenvalue image')

# show thresholded minor eigenvalue images
plt.subplot(3,2,3)
plt.imshow(J1_1, cmap='gray')
plt.title('thresholded minor eigenvalue image')
plt.subplot(3,2,4)
plt.imshow(J2_1, cmap='gray')
plt.title('thresholded minor eigenvalue image')

# show corners on original images
ax = plt.subplot(3,2,5)
plt.imshow(I1)
for i in range(C1): # draw rectangles of size w around corners

```

```

        x,y = pts1[:,i]
        ax.add_patch(patches.Rectangle((x-w/2,y-w/2),w,w, fill=False))
plt.title('found %d corners'%C1)

ax = plt.subplot(3,2,6)
plt.imshow(I2)
for i in range(C2):
    x,y = pts2[:,i]
    ax.add_patch(patches.Rectangle((x-w/2,y-w/2),w,w, fill=False))
plt.title('found %d corners'%C2)

plt.show()

```

Final values for parameters

- w =
- t =
- w_nms =
- C1 =
- C2 =

1.3 Problem 2 (Programming): Feature matching (15 points)

Determine the set of one-to-one putative feature correspondences by performing a brute-force search for the greatest correlation coefficient value (in the range $[-1, 1]$) between the detected features in image 1 and the detected features in image 2. Only allow matches that are above a specified correlation coefficient threshold value (note that calculating the correlation coefficient allows for adjusting the size of the matching window without changing the threshold value). Further, only allow matches that are above a specified distance ratio threshold value, where distance is measured to the next best match for a given feature. Vary these parameters such that around 200 putative feature correspondences are established. Optional: constrain the search to coordinates in image 2 that are within a proximity of the detected feature coordinates in image 1.

Report your final values for:

- the size of the matching window
- the correlation coefficient threshold
- the distance ratio threshold
- the size of the proximity window (if used)
- the resulting number of putative feature correspondences (i.e. matched features)

Display figures for:

- pair of images, where the matched features in each of the images are indicated by a square window about the feature and a line segment is drawn from the feature to the coordinates of the corresponding feature in the other image

A typical implementation takes around 10 seconds to run. If yours takes more than 30 seconds, you may lose points.

```

In [ ]: def NCC(I1, I2, pts1, pts2, w, p):
    # compute the normalized cross correlation between image patches I1, I2
    # result should be in the range [-1,1]
    #
    # inputs:
    # I1, I2 are the input images
    # pts1, pts2 are the point to be matched
    # w is the size of the matching window to compute correlation coefficients
    # p is the size of the proximity window
    #
    # output:
    # normalized cross correlation matrix of scores between all windows in
    # image 1 and all windows in image 2

    """your code here"""

    return scores

def Match(scores, t, d):
    # perform the one-to-one correspondence matching on the correlation coefficient matrix
    #
    # inputs:
    # scores is the NCC matrix
    # t is the correlation coefficient threshold
    # d distance ration threshold
    #
    # output:
    # 2xM array of the feature coordinates in image 1 and image 2,
    # where M is the number of matches.

    """your code here"""
    inds = []

    return inds

def RunFeatureMatching(I1, I2, pts1, pts2, w, t, d, p):
    # inputs:
    # I1, I2 are the input images
    # pts1, pts2 are the point to be matched
    # w is the size of the matching window to compute correlation coefficients
    # t is the correlation coefficient threshold
    # d distance ration threshold
    # p is the size of the proximity window
    #
    # outputs:

```

```

# inds is a 2xk matrix of matches where inds[0,i] indexes a point pts1
#     and inds[1,i] indexes a point in pts2, where k is the number of matches

scores = NCC(I1, I2, pts1, pts2, w)
inds = Match(scores, t, d, p)
return inds

In [ ]: # parameters to tune
w = 1
t = 1
d = 1
p = np.inf

tic = time.time()
# run the feature matching algorithm on the input images and detected features
inds = RunFeatureMatching(I1, I2, pts1, pts2, w, t, d, p)
toc = time.time() - tic

print('took %f secs'%toc)

# create new matrices of points which contain only the matched features
match1 = pts1[:,inds[0,:].astype('int')]
match2 = pts2[:,inds[1,:].astype('int')]

# display the results
plt.figure(figsize=(14,24))
ax1 = plt.subplot(1,2,1)
ax2 = plt.subplot(1,2,2)
ax1.imshow(I1)
ax2.imshow(I2)
plt.title('found %d putative matches'%match1.shape[1])
for i in range(match1.shape[1]):
    x1,y1 = match1[:,i]
    x2,y2 = match2[:,i]
    ax1.plot([x1, x2],[y1, y2], '-r')
    ax1.add_patch(patches.Rectangle((x1-w/2,y1-w/2),w,w, fill=False))
    ax2.plot([x2, x1],[y2, y1], '-r')
    ax2.add_patch(patches.Rectangle((x2-w/2,y2-w/2),w,w, fill=False))

plt.show()

# test 1-1
print('unique points in image 1: %d'%np.unique(inds[0,:]).shape[0])
print('unique points in image 2: %d'%np.unique(inds[1,:]).shape[0])

```

Final values for parameters

- w =

- $t =$
- $d =$
- $p =$
- $\text{num_matches} =$

1.4 Problem 3 (Programming): Outlier Rejection (15 points)

The resulting set of putative point correspondences should contain both inlier and outlier correspondences (i.e., false matches). Determine the set of inlier point correspondences using the M-estimator Sample Consensus (MSAC) algorithm, where the maximum number of attempts to find a consensus set is determined adaptively. For each trial, you must use the 4-point algorithm (as described in lecture) to estimate the planar projective transformation from the 2D points in image 1 to the 2D points in image 2. Calculate the (squared) Sampson error as a first order approximation to the geometric error.

hint: this problem has codimension 2

Report your values for:

- the probability p that at least one of the random samples does not contain any outliers
- the probability α that a given point is an inlier
- the resulting number of inliers
- the number of attempts to find the consensus set
- the tolerance for inliers
- the cost threshold

Display figures for:

- pair of images, where the inlier features in each of the images are indicated by a square window about the feature and a line segment is drawn from the feature to the coordinates of the corresponding feature in the other image

```
In [ ]: from scipy.stats import chi2
```

```
def DisplayResults(H, title):
    print(title+' =')
    print (H/np.linalg.norm(H)*np.sign(H[-1,-1]))

def Homogenize(x):
    # converts points from inhomogeneous to homogeneous coordinates
    return np.vstack((x,np.ones((1,x.shape[1]))))

def Dehomogenize(x):
    # converts points from homogeneous to inhomogeneous coordinates
    return x[:-1]/x[-1]

def MSAC(pts1, pts2, thresh, tol, p):
    # Inputs:
    # pts1 - matched feature correspondences in image 1
```

```

# pts2 - matched feature correspondences in image 2
# thresh - cost threshold
# tol - reprojection error tolerance
# p - probability that as least one of the random samples does not contain any outliers
#
# Output:
# consensus_min_cost - final cost from MSAC
# consensus_min_cost_model - planar projective transformation matrix H
# inliers - list of indices of the inliers corresponding to input data
# trials - number of attempts taken to find consensus set

"""your code here"""

trials = 0
max_trials = np.inf
consensus_min_cost = np.inf
consensus_min_cost_model = np.zeros((3,4))
inliers = np.random.randint(0, 200, size=100)
return consensus_min_cost, consensus_min_cost_model, inliers, trials

# MSAC parameters
thresh = 0
tol = 0
p = 0
alpha = 0

tic=time.time()

cost_MSAC, H_MSAC, inliers, trials = MSAC(pts1, pts2, thresh, tol, p)

# choose just the inliers
x1 = pts1[:,inliers]
x2 = pts2[:,inliers]

# Sample inliers. Please comment these before submission.
# xin1=np.loadtxt('hw4_pts1.txt').T
# xin2=np.loadtxt('hw4_pts2.txt').T

toc=time.time()
time_total=toc-tic

# display the results
print('took %f secs'%time_total)
print('%d iterations'%trials)
print('inlier count: ',len(inliers))

```

```

print('inliers: ',inliers)
print('MSAC Cost=%.9f'%cost_MSAC)
DisplayResults(H_MSAC, 'H_MSAC')

# display the figures
plt.figure(figsize=(14,24))
ax1 = plt.subplot(1,2,1)
ax2 = plt.subplot(1,2,2)
ax1.imshow(I1)
ax2.imshow(I2)
plt.title('found %d inliers'%xin1.shape[1])
for i in range(xin1.shape[1]):
    x1,y1 = xin1[:,i]
    x2,y2 = xin2[:,i]
    ax1.plot([x1, x2],[y1, y2],'-r')
    ax1.add_patch(patches.Rectangle((x1-w/2,y1-w/2),w,w, fill=False))
    ax2.plot([x2, x1],[y2, y1],'-r')
    ax2.add_patch(patches.Rectangle((x2-w/2,y2-w/2),w,w, fill=False))

plt.show()

```

Final values for parameters

- $p =$
- $\alpha =$
- tolerance =
- threshold =
- num_inliers =
- num_attempts =

1.5 Problem 4 (Programming): Linear Estimate (15 points)

Estimate the planar projective transformation \mathbf{H}_{DLT} from the resulting set of inlier correspondences using the direct linear transformation (DLT) algorithm (with data normalization). You must express $\mathbf{x}'_i = \mathbf{H}\mathbf{x}_i$ as $[\mathbf{x}'_i]^\perp \mathbf{H}\mathbf{x}_i = \mathbf{0}$ (not $\mathbf{x}'_i \times \mathbf{H}\mathbf{x}_i = \mathbf{0}$), where $[\mathbf{x}'_i]^\perp \mathbf{x}'_i = \mathbf{0}$, when forming the solution. Return \mathbf{H}_{DLT} , scaled such that $\|\mathbf{H}_{\text{DLT}}\|_{\text{Fro}} = 1$

```

In [ ]: def DLT(x1, x2, normalize=True):
    # Inputs:
    #   x1 - inhomogeneous inlier correspondences in image 1
    #   x2 - inhomogeneous inlier correspondences in image 1
    #   normalize - if True, apply data normalization to x1 and x2
    #
    # Outputs:
    #   H - the DLT estimate of the planar projective transformation
    #   cost - linear estimate cost

    """your code here"""

```

```

    # data normalization
    if normalize:
        print('normalize')

    # data denormalize
    if normalize:
        print('denormalize')

    H = np.eye(3)/np.sqrt(3) # estimated homography
    cost = np.inf # linear cost
    return H, cost

# compute the linear estimate without data normalization
print ('Running DLT without data normalization')
time_start=time.time()
H_DLT, cost = DLT(x1, x2, normalize=False)
time_total=time.time()-time_start
# display the results
print('took %f secs'%time_total)
print('Cost=%.9f'%cost)

# compute the linear estimate with data normalization
print ('Running DLT with data normalization')
time_start=time.time()
H_DLT, cost = DLT(x1, x2, normalize=True)
time_total=time.time()-time_start
# display the results
print('took %f secs'%time_total)
print('Cost=%.9f'%cost)

In [ ]: # display your H_DLT, scaled with its frobenius norm
        DisplayResults(H_DLT, 'H_DLT')

```

1.6 Problem 5 (Programming): Nonlinear Estimate (45 points)

Use H_{DLT} and the Sampson corrected points (in image 1) as an initial estimate to an iterative estimation method, specifically the sparse Levenberg-Marquardt algorithm, to determine the Maximum Likelihood estimate of the planar projective transformation that minimizes the reprojection error. You must parameterize the planar projective transformation matrix and the homogeneous 2D scene points that are being adjusted using the parameterization of homogeneous vectors (see section A6.9.2 (page 624) of the textbook, and the corrections and errata).

Report the initial cost (i.e. cost at iteration 0) and the cost at the end of each successive iteration. Show the numerical values for the final estimate of the planar projective transformation matrix

H_{LM} , scaled such that $\|H_{LM}\|_{Fro} = 1$.

```
In [ ]: from scipy.linalg import block_diag

def LM(H, x1, x2, max_iters, lam):
    # Input:
    #   H - DLT estimate of planar projective transformation matrix
    #   x1 - inhomogeneous inlier points in image 1
    #   x2 - inhomogeneous inlier points in image 2
    #   max_iters - maximum number of iterations
    #   lam - lambda parameter
    # Output:
    #   H - Final H (3x3) obtained after convergence

    # data normalization

    """your code here"""

    for i in range(max_iters):
        cost = np.inf
        print ('iter %03d Cost %.9f'%(i+1, cost))

    # data denormalization

    return H

# LM hyperparameters
lam = .001
max_iters = 100

# Run LM initialized by DLT estimate with data normalization
print ('Running sparse LM with data normalization')
print ('iter %03d Cost %.9f'%(0, cost))
time_start=time.time()
H_LM = LM(H_DLT, x1, x2, max_iters, lam)
time_total=time.time()-time_start
print('took %f secs'%time_total)

In [ ]: # display your converged H_LM, scaled with its frobenius norm
DisplayResults(H_LM, 'H_LM')
```